arXiv:1207.0558v4 [stat.AP] 25 Sep 2012

Bayesian semi-parametric forecasting of ultrafine particle number concentration with penalised splines and autoregressive errors

Sam Clifford[a,*], Bjarke Mølgaard[b], Sama Low Choy[c], Jukka Corander[d], Kaarle Hämeri[b], Kerrie Mengersen[c], Tareq Hussein[b,e]

[a]*School of Chemistry, Physics and Mechanical Engineering, Science and Engineering Faculty, Queensland University of Technology, GPO Box 2434, Brisbane, Qld 4000, Australia*
[b]*Department of Physics, University of Helsinki, P.O. Box 48, FI-00014, Finland*
[c]*School of Mathematical Sciences, Science and Engineering Faculty, Queensland University of Technology, GPO Box 2434, Brisbane, Qld 4000, Australia*
[d]*Department of Mathematics and Statistics, University of Helsinki, P.O. Box 68, FI-00014, Finland*
[e]*Department of Physics, University of Jordan, Amman, 11942, Jordan*

Abstract

Observational time series data often exhibit both cyclic temporal trends and autocorrelation and may also depend on covariates. As such, there is a need for flexible regression models that are able to capture these trends and model any residual autocorrelation simultaneously. Modelling the autocorrelation in the residuals leads to more realistic forecasts than an assumption of independence. In this paper we propose a method which combines spline-based semi-parametric regression modelling with the modelling of auto-regressive errors.

The method is applied to a simulated data set in order to show its efficacy and to ultrafine particle number concentration in Helsinki, Finland, to show its use in real world problems.

Keywords: autoregressive errors, semi-parametric regression, Bayesian inference, generalised additive model, ultrafine particles, aerosols

1. Introduction

Continuously measured time series may exhibit regular temporal trends, non-linear dependence on covariates (and interactions thereof) and autocorrelation. This motivates the desire for flexible regression models which are able to take these features into account without specifying the functional form of the relationship *a priori*.

The use of splines for non- and semi-parametric modelling smooth curves and surfaces (Silverman, 1985), time series (Wahba, 1990) and non-linear covariate effects (Lin and Zhang, 1999) is well established (Ruppert, Wand and Carroll, 2009). Splines have simple to construct bases (de Boor, 1978) and can be extended to include smoothness penalties, periodic bases, and interactions (Eilers and Marx, 2010). The flexibility of splines has led to their adoption as bases for Generalised Additive Models (Lin and Zhang, 1999; Lang and Brezger, 2004; Wood, 2006).

Harvey and Koopman (1993) describe a spline based model which flexibly models periodic trends by allowing the spline coefficients to evolve according to a random walk. This model admits the use of covariates but the error terms are assumed independent and identically distributed.

An important consideration when seeking to fit a model for forecasting time series data is that the residuals may display some degree of autocorrelation. Rather than fitting the mean predictor with an assumption of independent and identically distributed residuals and then performing *post hoc* analysis of the residual autocorrelation, including the autocorrelation in the modelling will ensure that the estimates of the model parameters have taken the autocorrelation into account, leading to more realistic forecasting (Chib, 1993; Mølgaard,

*Corresponding author
Email address: sj.clifford@student.qut.edu.au (Sam Clifford)

Hussein, Corander and Hämeri, 2012).

Traditional time series approaches (e.g. Box and Jenkins, 1994) are able to account for error structures other than independent and identically distributed errors through the specification of the Moving Average component of an ARIMA model (Venables and Ripley, 2002). Seasonality can be modelled in the ARIMA framework by including a term of the form $(1 - L^s)$, where L is the lag operator and s is the length of the period in terms of the frequency of the time series. The inclusion of this term accounts for seasonality by effectively removing the seasonal trend at lag s. Seasonal decomposition with loess (Cleveland, Cleveland, McRae and Terpenning, 1990; Cleveland, Grosse and Shyu, 1993) can perform interpolation and smooth estimation of seasonal trends with local polynomial regression. A common use of these seasonal decomposition models is to seasonally adjust a time series in order to examine the residual trends (Findley, Monsell, Bell, Otto and Chen, 1998), rather than examining the temporal trends themselves. Forecasting from these seasonally decomposed local regression models is possible, with ARIMA modelling providing the forecasts of the departures from the seasonal patterns.

One of the most flexible and easily implemented bases for the Generalised Additive Model is the penalised B-spline (Eilers and Marx, 1996). Despite the development of the GLM with autocorrelated errors (Chib, 1993) and the simplicity of the B-spline there does not appear to be an attempt made to incorporate both of these modelling approaches. A Bayesian nonparametric regression model with autocorrelated errors was proposed and implemented by Smith, Wong and Kohn (1998) which uses cubic regression splines rather than penalised splines. Penalised B-splines are attractive because when a large number of spline basis functions are used, any excessive wiggliness is penalised and a smoother fit obtained.

In this paper a method which combines spline-based semi-parametric regression modelling with the modelling of auto-regressive errors is proposed. It is possible to define a model for autoregressive (AR) errors in WinBUGS (Lunn, Thomas, Best and Spiegelhalter, 2000; Hay and Pettitt, 2001) but incorporating them with complex spline-based regression fixed effects terms is cumbersome in this modelling environment (Crainiceanu, Ruppert and Wand, 2005). The approach outlined below is a combination of the work of Mølgaard et al. (2012) and Clifford, Low Choy, Hussein, Mengersen and Morawska (2011).

Section 2 features a review of the Generalised Additive Model (GAM) and autocorrelated errors. The use of penalised splines is outlined, including how to construct penalties for multivariate bases and how to ensure the identifiability of the model (as the spline basis has full rank). Gibbs sampling for the spline coefficients and parameters for the autoregressive errors are discussed, as is the Metropolis sampler for the hyperprior which penalises the wiggliness of the splines. Forecasting with AR errors is also discussed, including posterior checks for predictive performance.

In Sections 3 and 4 the model is applied to two data sets. The first is a set of simulated time series data with a two dimensional (2D) covariate and autocorrelated errors added. This shows that the model recovers the temporal trend, autocorrelation of the errors and the 2D covariate effect. The posterior covariance of the sampled residuals are analysed to show the reduction in autocorrelation. The full regression and forecasting model is applied to observations of PNC and meteorology recorded in Helsinki.

In Section 5 the features of the proposed model are reviewed in light of the case studies and suggestions made for extensions to the modelling approach presented.

2. Methodology

Bayesian regression modelling can be conceptualised as starting with a set of assumptions about model parameters (prior belief) and using collected data to update those assumptions (to obtain the posterior belief). Mathematically, the prior beliefs about the parameters, $\boldsymbol{\theta}$, such as their mean and variance, are represented as a distribution, $p(\boldsymbol{\theta})$. The regression model and data are represented by a distribution where the likelihood of the data (X) is conditioned on the parameters, $p(\mathrm{X} \mid \boldsymbol{\theta})$. The posterior is obtained, through Bayes' rule by multiplying the prior and the likelihood, $p(\boldsymbol{\theta} \mid \mathrm{X})$, and represents a probability distribution for the parameters, conditioned on the observed data [1].

Chib (1993) specifies a Bayesian linear model in which the responses have either a Gaussian or t

[1] For a more thorough review of Bayesian statistics, see Gelman et al. (2004)

likelihood and the residuals are assumed to be autocorrelated rather than independent and identically distributed. In this section, an additive model with autocorrelated errors is developed by using splines to estimate non-linear covariate effects (Hastie and Tibshirani, 1990). The benefits of spline regression over polynomial regression have been discussed by Ruppert et al. (2003) and Crainiceanu et al. (2005).

Current software packages for Bayesian semi- and non-parametric Bayesian regression, such as R-INLA (Rue and Martino, 2010), mgcv[2] (Wood, 2011) and BayesX (Belitz et al., 2009), do not deal with autocorrelated residuals.

The methodology is implemented in MATLAB. The regression model is fit to the data which is observed and whose lagged data has also been observed. That is, we omit any data where any of y_i or its lagged values are missing (represented in MATLAB as `NaN`, "not a number").

2.1. Generalised Additive Models

Generalised Additive Models (GAMs) can be treated as equivalent to semiparametric regression with the Generalised Linear Mixed Model (GLMM) where the nonparametric smoothers, such as splines, are the random effects (Ruppert, Wand and Carroll, 2003; Fong, Rue and Wakefield, 2007). Alternatively, the basis matrix for the parametric part of the model can be augmented with the basis matrices for said smoothers, providing for the conception of the GAM as replacing the linear terms in the GLM with additive smoothers (Hastie and Tibshirani, 1990).

A GAM can therefore be expressed as

$$\mu_i = \beta_0 + f_1(x_{i1}) + f_2(x_{2i}) + \ldots + f_M(x_{Mi})$$
$$\boldsymbol{\mu} = \mathrm{X}\boldsymbol{\beta}$$

for M covariates, x_{m*}, each with their own univariate spline, $f_m(x_{m*})$, and a link function, $g(\cdot)$, as in the GLM. The design matrix of the GLM, X, now represents the spline design matrix. It has as many rows as there are observations and $\sum_{m=1}^{M} d_m$ columns, where d_m is the number of basis splines used in the construction of $f_m(\cdot)$.

The GAM can thus be written in a form which is equivalent to the formulation of a GLM, but with splines replacing the linear terms in the predictor,

$$g(\boldsymbol{y}) = \mathrm{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon}. \tag{1}$$

[2]while not strictly Bayesian, mgcv gives Bayesian summaries

2.2. Penalised splines

2.2.1. B-splines

The B-spline basis is built according to a recursive algorithm (de Boor, 1977; Eilers and Marx, 1996). A group of step functions, each of which is non-zero between two adjacent knots, t_{j-1} and t_j, are operated on to build a group of continuous linear functions which are non-zero between three adjacent knots such that the basis functions overlap. This process is repeated until the desired basis order is obtained.

The order k B-spline for covariate x

$$\frac{B_{j,k}(x)}{t_{j+k} - t_j} = \frac{x - t_j}{t_{j+k-1} - t_j}\frac{B_{j,k-1}(x)}{t_{j+k-1} - t_j} + \frac{t_{j+k} - x}{t_{j+k} - t_j}\frac{B_{j+1,k-1}(x)}{t_{j+k} - t_{j+1}} \tag{2}$$

on a grid of knots t with $B_{j,0}(x) = 1$ between knots t_{j-1} and t_j. That is, the zero-order B-spline basis is constant between successive knots and higher order bases are created according to the recurrence relation.

A cubic B-spline basis function consists of three quadratic polynomials which are piecewise continuous and smooth in the intervals between four neighbouring knots and the basis outside these knots is identically zero. Each B-spline basis function has compact support (i.e. are non-zero only on an open interval) and so allow the flexible modelling of non-linear effects by responding to local changes in the relationship between the corresponding covariate and the response.

The basis splines are calculated once for each non-linear function to be approximated and so the choice of the order of spline doesn't affect the complexity or speed of the calculations in the MCMC simulation, although the number of knots chosen will do so. The choice of a particular order of spline basis can therefore be made according to the desired continuity, smoothness and differentiability properties of the fitted smooth. A B-spline basis function will have non-zero derivatives up to and including the order of the basis spline, e.g. a first order spline will be once differentiable.

Figure 1 shows ten first order B-splines, ten second order B-splines and a linear combination of the same to approximate the function $y = \sin(2\pi x^2)$. The first order splines are piecewise continuous and that the second order splines are piecewise smooth. Each individual basis function is coloured a different shade of grey; the colouring scheme is consistent

across each of the four figures. The approximation is shown as a thick grey line and the true function is shown as a thick, black, dashed line.

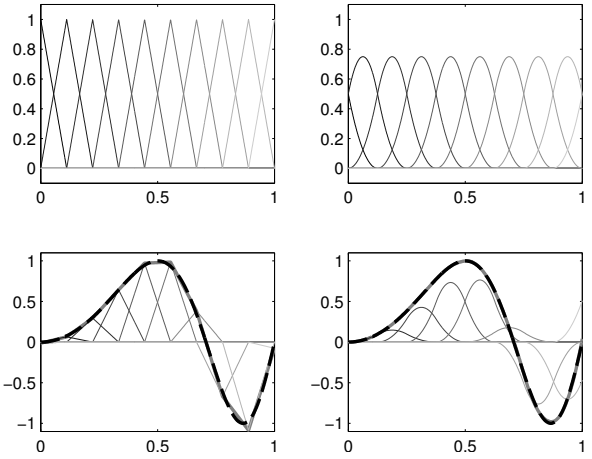

Figure 1: Example B-splines. Top row: ten first order B-splines, ten second order B-splines. Bottom row: Use of B-splines to approximate $y = \sin\left(2\pi x^2\right)$ where the true function is a dashed black line and the approximation is a solid grey line.

Figure 2 shows a 2D tensor basis of B-splines.

A cyclic B-spline basis can be defined by adding the basis splines with non-zero value at the edge of the univariate covariate space such that the basis is piecewise continuous at the boundary and any order preserving permutation of the knots on the periodic covariate space results in a valid B-spline basis.

2.2.2. Penalisation

A smoothing penalty can be imposed to ensure that the fitted B-splines are not too "wiggly". Eilers and Marx (1996) introduce the penalised B-spline, or "P-spline", in a frequentist context, modifying the normal equations to include a penalty matrix based on discrete differences between the spline coefficients.

Let $\mathrm{D} = \mathrm{D}(k)$ be a discretised operator equivalent to the kth derivative of the identity matrix with dimension equal to the number of coefficients in the B-spline basis of interest. This matrix will be rank deficient, with the deficiency in rank being equal to the order of the derivative. Applying the differential operator to a vector $\boldsymbol{\beta} \in \mathbb{R}^d$, where d is the number of basis elements in a univariate spline whose coefficients are $\boldsymbol{\beta}$, yields a vector with a value of the discretised value of that differential operator

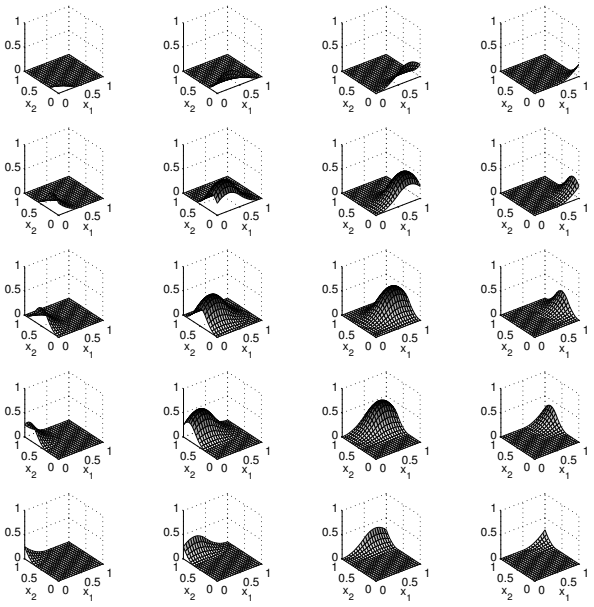

Figure 2: Tensor product of two B-splines. x_1 is modelled with four second order B-splines, x_2 with five second order B-splines; the tensor product thus has 20 basis elements. Shading is according to the value of the joint basis. Each image represents a 2D basis spline in the joint basis as a matrix, the spline basis matrix will consist of these matrices reshaped as column vectors.

evaluated at $\boldsymbol{\beta}$, and has a dimension equal to the number of rows of the differential operator, i.e. for $\boldsymbol{\beta} \in \mathbb{R}^d, \mathrm{D} \in \mathbb{R}^{(d-k)\times d}$, so that $\mathrm{D}\boldsymbol{\beta} \in \mathbb{R}^{(d-k)\times 1}$.

As $\mathrm{D}\boldsymbol{\beta}$ is the application of the discretised differential operator, $\mathbf{1}^T\mathrm{D}\boldsymbol{\beta} = \sum_{i=1}^{d-k} \mathrm{D}_{i,*}\boldsymbol{\beta}$, is the sum of the discretised derivatives. The continuous analogy is $\int f^{(k)}(x)\,\mathrm{d}x$. Then

$$\left(\boldsymbol{\beta}^T\mathrm{D}^T\right)(\mathrm{D}\boldsymbol{\beta}) = \int \left[f^{(k)}\right]^2 \,\mathrm{d}x$$

and is the discretised sum of the squares of the derivatives (Wood, 2003). The quantity to be minimised in the penalised B-spline, then, is the sum of the square of the derivative (Eilers and Marx, 1996).

Lang and Brezger (2004) reformulate this penalty matrix as a zero mean conjugate multivariate normal prior on the B-spline coefficients, $\boldsymbol{\beta} \in \mathbb{R}^{d\times 1}$,

$$(\boldsymbol{\beta} \mid \lambda) \propto \lambda^{(d-k)/2} \exp\left(-\frac{\lambda}{2}\boldsymbol{\beta}^T\mathrm{K}\boldsymbol{\beta}\right) \tag{3}$$

$$\lambda \sim \Gamma(1, b) \tag{4}$$

where λ is the smoothing parameter which penalises the "wiggliness" of the resulting spline and $\mathrm{K} \in \mathbb{R}^{d \times d}$ is the matrix form of the discretised wiggliness penalty as in Eilers and Marx (1996). The parameter for the scale of the prior on λ, b, is chosen such that the prior has a small variance and is concentrated around $\lambda = 0$. The prior for $\boldsymbol{\beta}$ is improper, but a non-diffuse prior on λ ensures that the posterior for $\boldsymbol{\beta}$ is proper. The rank deficiency of the penalty matrix is overcome by adding a small amount to the diagonal of K of no more than about 0.001 (the diagonal values of K are at least 2, the first order penalty case).

There are other ways to penalise the B-spline coefficients than the P-spline approach of Eilers and Marx (2003). These include the approach of O'Sullivan (1988) and its Bayesian analogue, developed by Wand and Ormerod (2008).

2.2.3. Interaction terms

Harvey and Koopman (1993) describe a spline based model which is able to flexibly model periodic trends by allowing the spline coefficients to evolve according to a random walk. This model is applied to a joint model of daily and hourly trends and treats the coefficients of the daily trend spline as evolving according to a random walk model from day to day. While this is an obvious choice for modelling an interaction between a factor and a continuous covariate, it does not make sense for two continuous covariates.

A tensor product of the bases for two univariate B-spline matrices can be constructed in order to fit an interaction term of those covariates (Eilers and Marx, 2003; Marx and Eilers, 2005). Penalisation of the coefficients for this tensor basis is achieved by forming the Kronecker tensor product, denoted $\otimes$, of the identity matrix with size equal to the number of basis elements and the univariate penalty matrix for each univariate spline. This leads to a formulation of the smoothing penalty prior in terms of two penalty matrices, one along each axis of the covariate space corresponding to each of the two covariates. The prior takes the form

$$\begin{aligned}(\boldsymbol{\beta} \mid \lambda_1, \lambda_2) &\propto \exp\left(-\frac{\lambda_1}{2}\boldsymbol{\beta}^T \mathrm{P}_1 \boldsymbol{\beta} - \frac{\lambda_2}{2}\boldsymbol{\beta}^T \mathrm{P}_2 \boldsymbol{\beta}\right) \\ \lambda_1, \lambda_2 &\sim \Gamma(1, b) \\ \mathrm{P}_1 &= K_1 \otimes \mathrm{I}_{d_2} \\ \mathrm{P}_2 &= \mathrm{I}_{d_1} \otimes K_2\end{aligned} \tag{5}$$

for d_1, d_2 the number of basis elements for covariates 1 and 2 respectively. The matrices P_1 and P_2 are square matrices of dimension $d_1 \times d_2$, the first of which is block diagonal, with each block composed of tridiagonal matrices, corresponding to the penalties on the first covariate's basis (indexed sequentially). The second, P_2, has a banded structure, corresponding to the penalties on the second covariate's basis (where the indexing is sequential, d_1 at a time). An example of a 2D tensor product of B-splines is shown in Figure 2 in the appendix.

The above tensor product formulation with a penalty for each univariate basis reflects the possibility that a different amount of smoothing may apply in each direction, i.e. that smoothing is not isotropic. While R-INLA provides a `rw2d` GMRF latent model for two-dimensional smooth terms, the single precision parameter of the GMRF (equivalent in formulation to a univariate P-spline prior) does not reflect that there may be a different amount of smoothing required in each direction. This multiple penalty is addressed in BayesX, along with an interaction between the smoothing terms (Belitz et al., 2009).

The random walk time evolving spline of Harvey and Koopman (1993) can be recovered by using a first order random walk P-spline penalty matrix as the precision matrix for a factor term.

2.3. Parametric terms

In this section, the term "parametric" denotes regression basis functions with a parametric form specified *a priori*. These parametric terms, such as linear terms, polynomials and sinusoidal functions, will be modelled as fixed effects with Gaussian priors (multivariate where appropriate), according to the method outlined by Chib (1993).

In section 2.2.1 the periodic B-spline basis was briefly discussed. Another choice of basis for semi-parametric modelling of periodic functions is the Fourier series basis, consisting of $\sin\left(\frac{c\pi x}{L}\right)$, $\cos\left(\frac{c\pi x}{L}\right)$, $c = 1 \ldots C$. The full Fourier series basis also includes an intercept term which should be omitted from the regression model to ensure identifiability of the overall mean. In the case of the splines section 2.6.1 discusses how identifiability is enforced by centring the splines; the sin and cos terms are already centred.

The sinusoidal functions may be treated as any other parametric (e.g. linear) term in the regression equation. It is appropriate to assume that the

Fourier coefficients are uncorrelated since by definition the Fourier basis is orthogonal. An appropriate prior, therefore, is a weakly informative multivariate normal with zero mean and a covariance matrix whose structure is diagonal.

Heiler and Feng (2000) provide an example of using the Fourier series basis to approximate temporal trends. As with all variable selection problems, the question is how many terms must be included? A difficulty of using Fourier series is that using too few terms constrains the "wiggliness" of the resulting approximated function whereas using terms with too high frequency may lead to spurious oscillations unless the prior shrinks the coefficients of higher frequency oscillations to zero as in Lenk (1999). Instead of this approach, the use of periodic B-splines with a smoothing penalty is adopted. Here it is merely pointed out that the Fourier basis may be used, as was done by Mølgaard et al. (2012).

2.4. Factor terms

Factor terms may arise as random effects in a GAMM, e.g. a site-specific mean in multi-site data, and may be treated by coding all levels of the factors to an integer, as in R-INLA. For a factor basis matrix with sequential integers indexing the factors, $\mathrm{F} \in \mathbb{R}^{n\times J}$, for covariate $\boldsymbol{x} \in \mathbb{R}^n$, the element $F_{i,j} = 1$ if $x_i = j$ $(j \in 1, 2, \ldots, J)$.

For the coprecision of the factor term, Q_F, it is simple to assume that the factors are independent and identically distributed, so that $\mathrm{Q}_F = \lambda_F \mathrm{I}$ for precision parameter λ_F. If there is some known structure in the factors or if a factor is being used to describe a zero order spline (constant between knots with knots spaced between factor indices) then a smoothing penalty prior may be used, as for a P-spline.

The factor terms are centred around zero with the same identifiability constraint as in 2.6.1.

2.5. Autocorrelated residuals

Many time series are highly autocorrelated. Any temporal variation not explained by a periodic spline basis (e.g. hour of day, day of year) in the proposed model can treated by modelling the time series of errors, ε_i, with an autoregression model,

$$\begin{aligned} \phi(L)\varepsilon_i &= u_i \\ u_i &\sim \mathcal{N}\left(0, \sigma^2\right). \end{aligned} \tag{6}$$

where L is the lag operator and $\phi(L)$ is a p-degree polynomial in the lag operator

$$\phi(L)x_i = x_i + \phi_1 x_{i-1} + \phi_2 x_{i-2} + \ldots + \phi_p x_{i-p}$$

describing the autocorrelation structure of the residuals in 1 (Box and Jenkins, 1994).

Chib (1993) suggests the use of a multivariate normal prior for the coefficients of the above polynomial. A censoring condition, that the absolute sum of the coefficients is strictly less than one, is applied in order to ensure that the error process is stationary. The prior is

$$\boldsymbol{\phi} \propto \mathcal{MVN}\left(\boldsymbol{\phi}_0, \Phi_0^{-1}\right) \ : \ \left|\mathbf{1}^T\boldsymbol{\phi}\right| < 1.$$

where the prior mean, $\boldsymbol{\phi}_0$, and precision, $\Phi_0 \in \mathbb{R}^{p\times p}$, are set such that the prior for the autoregression parameters is weakly informative. The prior for the variance of the autoregressive errors is a conjugate inverse Gamma,

$$\sigma^2 \sim \mathcal{IG}\left(\frac{v_0}{2}, \frac{\delta_0}{2}\right).$$

By using an autoregressive error structure, forecasting from the model will give a more realistic representation of the relationship between successive observations/forecasts. If the residuals are independent, this will manifest as the zero vector being contained within the multivariate credible set of $\boldsymbol{\phi}$.

By sampling the residuals in the MCMC scheme and using these for forecasting, a distribution for each forecast is readily obtainable. This will be especially important when making a forecast which is dependent on a previously forecast observation. Ignoring the uncertainty in the distribution of the forecast value $\widehat{y}_{t+1}$ would produce an estimate of $\widehat{y}_{t+2}$ which assumed that $\hat{y}_{t+1}$ was observed rather than forecast.

2.6. Estimation via MCMC

The form of the GAM with log link, Gaussian likelihood and autocorrelated residuals is

$$\begin{aligned} \log y_i &= \mu_i + \varepsilon_i \\ \mu_i &= \mathrm{X}_{*,i}\boldsymbol{\beta} \\ \phi(L)\varepsilon_i &= u_i \\ u_i &\sim \mathcal{N}\left(0, \sigma^2\right) \end{aligned} \tag{7}$$

where $\mathrm{X}_{i,*}$ is row i of matrix X.

A model for a Gaussian hierarchical linear model with autoregressive errors is given by Mølgaard et al. (2012), following the formulation of Chib (1993). The priors for σ^2, $\boldsymbol{\beta}$ and $\boldsymbol{\phi}$, in this formulation are

$$\begin{aligned}\left(\boldsymbol{\beta} \mid \sigma^2\right) &\sim \mathcal{MVN}\left(\overline{\boldsymbol{\beta}}_0, \sigma^2 \mathrm{A}_0^{-1}\right) \\ \sigma^2 &\sim \mathcal{IG}\left(\frac{v_0}{2}, \frac{\delta_0}{2}\right) \\ \boldsymbol{\phi} &\propto \mathcal{MVN}\left(\boldsymbol{\phi}_0, \Phi_0^{-1}\right) \ : \ \left|\mathbf{1}^T \boldsymbol{\phi}\right| < 1.\end{aligned} \tag{8}$$

The terms ϕ_0 and $\overline{\boldsymbol{\beta}}_0$ are the respective mean values for the weakly informative priors for $\boldsymbol{\phi}$ and $\boldsymbol{\beta}$, taken to be $\mathbf{0}$ in each case. The values for v_0 and δ_0 are taken to be $-k$ and 0, respectively, as in Chib (1993). The prior precision matrix for the autoregressive error model, Φ_0, is set as $10^{-6}\mathrm{I}_p$ so that the prior is weakly informative, such. The prior precision matrix for $\boldsymbol{\beta}$, A_0 is discussed below.

Chib sets the prior precision of $\boldsymbol{\beta}$, A_0, as a scalar multiple of an identity matrix such that the prior is weakly informative. Here, however, A_0 is a block diagonal with each block containing a smoothness penalty precision matrix, $\lambda_j \mathrm{K}_j$, for a spline term. The λ are therefore hyperparameters for the prior on $\boldsymbol{\beta}$. Any parametric terms to be included in the model can be treated similarly, by including their prior precision as a series of individual elements on the diagonal (in the case that they are independent), a block consisting of a diagonal matrix (for sets of independent, identically distributed parameters) or a block structured as a symmetric positive definite matrix (if there is some known correlation structure). The prior for $\boldsymbol{\beta}$ is now is very informative due to the spline penalties contained within A_0.

The posterior distributions for the linear predictor terms, $\boldsymbol{\beta}$, autoregressive coefficients, $\boldsymbol{\phi}$, and model variance, σ^2 are given in Table 1.

The matrix X^*, and the vector $\boldsymbol{y}^*$ are formed by applying the general lag polynomial $1 - \phi(L)$ to X and $\boldsymbol{y}$ respectively, reducing the number of rows by p, the number of lags included in the autoregressive residual term. That is,

$$\begin{aligned}\mathrm{X}^* = \mathrm{X}_{1\ldots n-p,*} - \phi_1 \mathrm{X}_{2\ldots n-p+1,*} - \\ \phi_2 \mathrm{X}_{3\ldots n-p+2,*} \ldots - \phi_p \mathrm{X}_{p+1\ldots n,*}\end{aligned} \tag{12}$$

and similarly for $\boldsymbol{y}^*$.

$\mathrm{E} \in \mathbb{R}^{(n-p)\times p}$ contains, as its columns, the residuals, $\boldsymbol{\varepsilon}$, at each of the p lags, such that the i^{th} row is $[\varepsilon_{i-1}, \ldots, \varepsilon_{i-p}]$,

$$\mathrm{E} = \left[\boldsymbol{\varepsilon} \mid L\boldsymbol{\varepsilon} \mid L^2\boldsymbol{\varepsilon} \mid \cdots \mid L^p\boldsymbol{\varepsilon}\right]. \tag{13}$$

In 10, $\|\cdot\|_2$ represents the Euclidean 2-norm.

Rather than first fitting the linear predictor with the assumption of independent, identically distributed errors and then fitting an autoregressive model for the residuals so that $\boldsymbol{\phi}$ is conditioned on $\boldsymbol{\beta}$ but not vice versa, this comprehensive model allows simultaneous estimates of $\boldsymbol{\beta}$ and $\boldsymbol{\phi}$.

The directed acyclic graph of the model is given in Figure 3. Sampling is performed by successively sampling from the marginals of $\boldsymbol{\beta}^{(t)}$ from 9 (then correcting $\boldsymbol{\beta}^{(t)}$ according to 14 and 15), $\boldsymbol{\phi}^{(t)}$ from 11, $\sigma^{2(t)}$ from 10 and sample $\boldsymbol{\lambda}^{(t)}$ from 17 and 18 with a Metropolis-Hastings step.

2.6.1. Identifiability of spline coefficients

A d-dimensional spline basis B will have full rank and therefore the matrix $[\mathbf{1}|\mathrm{B}] \in \mathbb{R}^{n\times(d+1)}$ is rank deficient. In order to enforce identifiability the B-spline parameters are centred at each step of the Gibbs sampler such that the sum of the marginal effect of that spline is zero (Wood, 2003) and centering constant is transferred to the intercept term in the model, β_0.

At step t in the Gibbs sampler, for each block of the covariate matrix $\mathrm{X} \in \mathbb{R}^{N\times d}$ corresponding to a B-spline, indexed within X by ind we set

$$\begin{aligned}\boldsymbol{\beta}_{\text{ind}}^{(t)} &\leftarrow \boldsymbol{\beta}_{\text{ind}}^{(t)} - \delta_\beta \mathbf{1} \\ \beta_0^{(t)} &\leftarrow \beta_0^{(t)} + \delta_\beta.\end{aligned} \tag{14}$$

where the centring constant, δ_β, is calculated as that which satisfies

$$\begin{aligned}0 &= \sum_{i=1}^{N} \left(\mathrm{X}_{i,\text{ind}} \boldsymbol{\beta}_{\text{ind}}^{(t)} - \mathrm{X}_{i,\text{ind}} \mathbf{1} \delta_\beta\right) \\ \delta_\beta &= \frac{\sum \mathrm{X}_{*,\text{ind}} \boldsymbol{\beta}_{\text{ind}}^{(t)}}{\sum_{i=1}^{N} \sum_{j\in\text{ind}} \mathrm{X}_{i,j}}.\end{aligned} \tag{15}$$

2.6.2. Lack of support

It is possible that B-spline knots will be placed in a region of the covariate space where there are no observed values, especially when dealing with interaction terms. For a basis spline defined in a region of no support (i.e. there are no observations in the relevant region of the covariate space) the

$$(\boldsymbol{\beta} \mid \boldsymbol{y}, \sigma^2, \boldsymbol{\phi}) \sim \mathcal{MVN}\left(\Lambda_\beta \left(\mathrm{A}_0\overline{\boldsymbol{\beta}}_0 + \mathrm{X}^{*T}\mathrm{y}^*\right), \sigma^2\Lambda_\beta\right) \tag{9}$$

$$(\sigma^2 \mid \boldsymbol{y}, \boldsymbol{\beta}, \boldsymbol{\phi}) \sim \Gamma^{-1}\left(\frac{n-p+v_0+k}{2}, \frac{\delta_0 + \mathrm{Q}_\beta + \|\mathrm{y}^* - \mathrm{X}^*\boldsymbol{\beta}\|_2^2}{2}\right) \tag{10}$$

$$(\boldsymbol{\phi} \mid \boldsymbol{y}, \boldsymbol{\beta}, \sigma^2) \sim \mathcal{MVN}\left(\Lambda_\phi \left(\Phi_0\phi_0 + \sigma^{-2}\mathrm{E}^T\boldsymbol{\varepsilon}\right), \Lambda_\phi\right) \tag{11}$$

where

$$\begin{aligned}
\Lambda_\beta &= \left(\mathrm{A}_0 + \mathrm{X}^{*T}\mathrm{X}^*\right)^{-1} \\
\mathrm{Q}_\beta &= \left(\boldsymbol{\beta} - \overline{\boldsymbol{\beta}}_0\right)^T \mathrm{A}_0 \left(\boldsymbol{\beta} - \overline{\boldsymbol{\beta}}_0\right) \\
\Lambda_\phi &= \left(\Phi_0 + \sigma^{-2}\mathrm{E}^T\mathrm{E}\right)^{-1}.
\end{aligned}$$

Table 1: Posterior densities for GAM with AR residuals

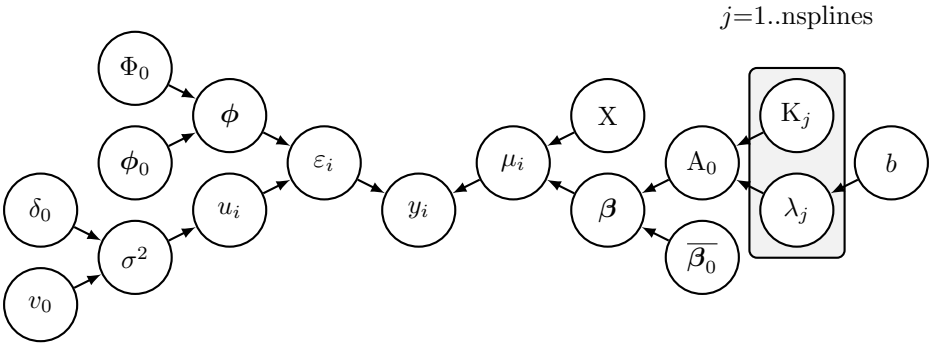

Figure 3: Directed acyclic graph of the model in 7 with priors 8, 3, 4 and 5.

posterior of the corresponding coefficient is undefined. In these cases, a sample will be taken from the prior. Due to the informative nature of the penalty prior, the posterior variance of this parameter will be much smaller than if the spline was unpenalised. The corresponding marginal effect of the entire spline will have a wider credible interval in the region of this spline's support.

An alternative is to use a spline basis with global support for the covariate responsible for the lack of support, such as the semi-parametric low rank thin plate spline described by Ruppert et al. (2003) and Crainiceanu et al. (2005) as

$$f(x_i) = \sum_{j=1}^{J} \beta_j x_i^j + \sum_{k=1}^{K} \gamma_k \left|x_i - \kappa_k\right|^m$$

where $J < m$ is the order of the parametric fixed effect polynomial, m is the order of the random effects polynomial splines and the κ_k are knots placed at the $k/(K+1)$ quantiles of x. A suitable prior for $(\beta_1, \ldots, \beta_J, \gamma_1, \ldots \gamma_K)$ is a weakly informative multivariate normal with an identity precision matrix.

2.6.3. Sampling spline penalties

In order to sample the P-spline penalty correctly we need to extend the Gibbs sampler to also sample the λ, which are hyperparameters for the prior distribution for $\boldsymbol{\beta}$. The penalty matrices K and P, suggested by Lang and Brezger (2004), however, have a determinant of zero and thus the posterior of λ is undefined. We avoid this by adding a small value to the diagonals of these matrices to ensure that they are non-singular and thus invertible, e.g. the penalty becomes $\lambda\left(\mathrm{K} + 10^{-5}\mathrm{I}\right)$.

For the one-dimensional splines it is rather simple

to find the conditional distribution of the λ hyperparameters when $\mathrm{K} \in \mathbb{R}^{d \times d}$ has non-zero determinant. The posterior is derived as

$$\begin{aligned} p(\lambda \mid \boldsymbol{\beta}) &\propto p(\boldsymbol{\beta} \mid \lambda)\, p(\lambda) \qquad (16) \\ &= \mathcal{MVN}\left(\boldsymbol{\beta} \,\middle|\, \mathbf{0}, (\lambda \mathrm{K})^{-1}\right) \Gamma(\lambda \mid 1, b) \\ &= \sqrt{\frac{\det(\lambda \mathrm{K})}{(2\pi)^d}} \exp\left(-\frac{\lambda}{2} \boldsymbol{\beta}^T \mathrm{K} \boldsymbol{\beta}\right) b e^{-b\lambda} \\ &\propto \lambda^{d/2} \exp\left(-\lambda\left(b + \frac{1}{2}\boldsymbol{\beta}^T \mathrm{K} \boldsymbol{\beta}\right)\right) \\ &\propto \Gamma\left(\lambda \,\middle|\, \frac{d}{2} + 1, b + \frac{1}{2}\boldsymbol{\beta}^T \mathrm{K} \boldsymbol{\beta}\right). \qquad (17) \end{aligned}$$

That is, the Gamma prior, 4, is updated by the data in such a way that some smoothing is enforced by the dimension of the spline but the scale parameter, which controls the amount of smoothing (or how likely the smoothing is), is shifted according to the second difference of the parameters. It is still possible to have zero smoothing, but the posterior density is small, so that a lack of smoothing must be justified by the data.

For two dimensional splines the sampling of λ would be as simple as the univariate case if the smoothing was set *a priori* to be the same for each covariate in the tensor product. It is more valid, though, to assume that this is not the case and that there may be more smoothness in one direction than the other, implying two different λ parameters. The conditional distribution of these parameters, assuming that the same hyperpriors for the Gamma prior are used for each spline, is

$$\begin{aligned} p(\lambda_1, \lambda_2 \mid \boldsymbol{\beta}) \propto \sqrt{\det(\lambda_1 \mathrm{P}_1 + \lambda_2 \mathrm{P}_2)}\, (\lambda_1 \lambda_2)^{a-1} \times \\ e^{-\frac{1}{2}\boldsymbol{\beta}^T(\lambda_1 \mathrm{P}_1 + \lambda_2 \mathrm{P}_2)\boldsymbol{\beta}} e^{-b(\lambda_1 + \lambda_2)}. \qquad (18) \end{aligned}$$

Because the conditional distributions of λ_1 and λ_2 are not easy to obtain, we sample them with a Metropolis-Hastings step (Hastings, 1970) inside the Gibbs sampler.

2.6.4. Credible intervals

Credible intervals of parameter estimates can be calculated quickly from the MCMC samples. For any parametric terms, 95% credible intervals will be reported. For the coefficients of splines, the credible interval for the entire marginal effect is of interest and so individual credible intervals for each parameter are not so directly interpretable for the spline terms.

2.7. Forecasting

The model forecasts in a sequential manner, conditioning each forecast on the observed values and forecast values preceding it according to the autoregressive structure. The uncertainty in forecast values will accumulate, as discussed in Section 2.5.

In Section 3.2, to take advantage of the growing set of training data the model is refitted every 20th day. For each day, forecasts are produced one time hour at a time for the coming forty-eight hours, midnight to midnight, using the most recent fit of the training data and the residuals, $\boldsymbol{\varepsilon}$, for the past week. These residuals are assumed to be available until noon. The forecast values for each observation made 24 and 48 hours in advance will be stored.

The modelled value of each forecast is obtained by calculating $\widehat{\boldsymbol{\mu}} = \mathrm{X}\widehat{\boldsymbol{\beta}}$ for the values to be forecast, based on the observed covariates (in practice, these may be known and the forecasting treated as imputation, they may be forecasts themselves from a model for the covariates, may be generated from a prior, etc.); calculating the residuals, $\widehat{\boldsymbol{\varepsilon}}$, for the week before the forecasting is to start; and constructing autocorrelated forecasting errors from the estimates of $\boldsymbol{\phi}$. These are all constructed from MCMC output, so distributions of these parameters are all available. The forecasts, $\widehat{y}_i$, are calculated from the MCMC iterations so that

$$\widehat{y}_i^{(t)} = \widehat{\mu}_i^{(t)} + \widehat{\varepsilon}_i^{(t)}.$$

Because both $\widehat{\mu}_i$ and $\widehat{\varepsilon}_i$ are normally distributed, $\widehat{y}_i$ is also normally distributed. The posterior density of the forecast values, then, is the empirical density of the samples and is asymptotically normal as the number of samples increases.

For further information, see Sections 2.1.3 to 2.1.5 of Mølgaard et al. (2012).

2.8. Posterior checks

Convergence of the MCMC chains can be checked by examining the trace of the Markov chains and that the posterior kernel density estimate is unimodal and that its shape is appropriate to the distribution from which it was drawn.

As the posterior for a fitted smooth is a multivariate normal it is important to remember that even though the 95% CI for a B-spline coefficient (which is marginally normally distributed) may contain zero, it is the joint effect that we are interested in. Indeed, because the spline is centred around

zero, it is not unexpected that zero would be contained in the credible interval of some coefficients. It is more informative to look at whether zero is contained within the credible interval for the entire smooth and therefore the spline can be said to be identically zero.

Plots of the posteriors of the marginal effects are a very informative way of both visualising splines and checking that the output of the model is consistent with the physical system that the statistical model is attempting to represent.

For plotting, we calculate a new basis with evenly spaced covariate values between the extreme values of each covariate in the data. This gives a basis with the same knots as the basis used for fitting the model but with a set of covariates that will give a visually more pleasing plot. The mean and 95% credible regions for these posteriors are calculated by using the output from the Gibbs sampler. Bivariate splines will be plotted as surface and/or contour plots.

For a Generalised Linear Mixed Model, which encompasses a GAM with splines for random effects (Fong et al., 2007),

$$\boldsymbol{y} = g^{-1}\left(\widetilde{\mathrm{X}}\widetilde{\boldsymbol{\beta}} + \widetilde{\mathrm{Z}}\widetilde{\boldsymbol{\gamma}}\right) + \boldsymbol{\varepsilon},$$

the fitted values are given by

$$\begin{aligned}\widehat{\boldsymbol{y}} =& \widetilde{\mathrm{X}}\widehat{\boldsymbol{\beta}} + \widetilde{\mathrm{Z}}\widehat{\boldsymbol{\gamma}} \\ =& \mathrm{X}\left(\mathrm{X}^T\mathrm{X} + \Lambda\right)^{-1}\mathrm{X}^T\boldsymbol{y}\end{aligned}$$

where $\mathrm{X} = \left[\widetilde{\mathrm{X}} \,\middle|\, \widetilde{\mathrm{Z}}\right] \in \mathbb{R}^{n\times(l+k)}$ (for l parametric terms and k non-parametric basis functions) and

$$\Lambda = \begin{bmatrix} 0_{l\times l} & 0_{l\times k} \\ 0_{k\times l} & \sigma^2\left(\mathrm{cov}(\boldsymbol{\gamma})\right)^{-1} \end{bmatrix} \tag{19}$$

is the posterior precision of $\boldsymbol{\beta}$ from 9 (Ruppert et al., 2003).

The total effective number of degrees of freedom of the model is given by the trace of the hat matrix

$$\mathrm{H} = \left(\mathrm{X}^T\mathrm{X} + \Lambda\right)^{-1}\mathrm{X}^T\mathrm{X} \tag{20}$$

and for each non-parametric term, the trace of the block matrix of the hat matrix corresponding to the appropriate coefficients gives the effective number of degrees of freedom for that non-parametric term (Fong et al., 2007).

Rather than estimate the mean effective degrees of freedom *post hoc* with the posterior means of the parameters we can use the samples from the MCMC chains as they are drawn to construct H. In this way, we obtain the distribution of the number of parameters for each term by examining the density estimate of the trace of the blocks of H.

The Deviance Information Criterion (DIC) is a Bayesian analogue of the Akaike Information Criterion and penalises the deviance, a measure of the goodness of fit of a model, by the effective number of parameters in the model (Spiegelhalter et al., 2002). The effective number of parameters can be compared to the effective degrees of freedom (described above). The effective number of parameters is based on the deviance and so the DIC does not require the integrating out of random effects parameters like Schwarz's Bayesian Criterion (Schwarz, 1978).

While calculation of the credible interval for $\boldsymbol{\phi}$ will provide information about the degree of autocorrelation of the residuals, the autocorrelation function of both $\boldsymbol{\varepsilon}$ and $\boldsymbol{u}$ can be used to check how much autocorrelation remains. The posterior covariance matrix from the chains of $\boldsymbol{\varepsilon}$ and $\boldsymbol{u}$ can also be used to characterise the remaining autocorrelation of the residuals.

The probability integral transform (PIT) will be used to assess the quality of the model fit. The PIT is a measure of the cumulative density of the forecast values from their predictive distribution (Dawid, 1984; Diebold, Gunther and Tay, 1998). This predictive CDF, denoted F_i, is to be compared to an unknown "true" CDF G_i through the observed values that the true physical process generates, y_i, and the forecast values, $\widehat{y}_i$. The ideal forecasting model is achieved when $F_i = G_i$. A necessary condition for choosing a forecasting model is that the distribution of $F_i(\widehat{y}_i)$ is uniform, corresponding to ideal forecasts.

The PIT for each forecast value is calculated during the forecasting step as

$$F_i(\widehat{y}_i) = \mathrm{ecdf}\left(\widehat{y}_i;\, \widehat{\mu}_i + \widehat{\varepsilon}_i\right)$$

where ecdf is the empirical cumulative density function of the posterior samples of $\widehat{\mu}_i + \widehat{\varepsilon}_i$ from the forecasting routine. The PIT is visualised by plotting the histogram and autocorrelation of $F_i(\widehat{y}_i)$. If the PIT is uniformly distributed, then the values of the probability density function corresponding to the PIT's cumulative density function should be a normal distribution with mean 0 and standard deviation 1. The autocorrelation of the PIT should

not be significant beyond the number of steps used in forecasting.

Many competing models may have a uniform PIT so the uniformity is not a sufficient condition for choosing one model over another (Hamill, 2001). To remedy this, Gneiting et al. (2007) recommend maximising the sharpness of the predictive distribution, i.e. choosing forecasts which are highly concentrated about the observed values. The variance of the forecast values will be estimated by characterising their uncertainty with the width of the 95% credible interval of the simulated forecasts (which are distributed normally). The 95% credible interval corresponds to 1.96 standard deviations of the estimate, so dividing the interval half width by 1.96 and taking the square will give an estimate of $\mathrm{Var}\left(\widehat{\boldsymbol{y}}\right)$.

3. Case studies

To demonstrate the performance of the methodology in terms of its ability to fit smooth covariate effects and autocorrelated residuals we provide two case studies. The first is a simulation study with a univariate cyclic smooth and a bivariate interaction term. The second case study involves modelling and forecasting ultrafine particle number concentration in Helsinki, Finland, a real world data set that exhibits a significant amount of autocorrelation.

3.1. Simulation

As a first case study, to illustrate the method, we provide an example with simulated data using a Gaussian likelihood, as in Chib (1993). Data is simulated as described in 21. Two covariates, x and y, are each drawn from a uniform distribution and a non-linear interaction term constructed that requires that a 2D spline be fit. A sinusoidal function to simulate a temporal trend (at evenly spaced intervals) is added, as is some autocorrelated noise.

$$\begin{aligned} \boldsymbol{x}, \boldsymbol{y} &\sim \mathcal{U}\left(0,1\right) \\ \boldsymbol{t} &= 1, 2, \ldots, 24 \\ \mu_i &= \sin\left(\pi x_i\right)\left(1 - x_i y_i^2\right) + \frac{\sin\left(\frac{2\pi t_i}{24}\right)}{2} \\ \left(1 + 0.4L\right)\varepsilon_i &\sim \mathcal{N}\left(0, 0.1\right) \end{aligned} \tag{21}$$

The model for this simulation thus contains: a univariate spline for $\boldsymbol{t}$ with a basis of six second order cyclic B-splines with a second order penalty prior; a bivariate spline for fitting $\boldsymbol{x}$ and $\boldsymbol{y}$ together, consisting of the tensor product of two second order non-cyclic B-splines bases with six basis splines each and a second order penalty prior for each direction; an intercept term with a weakly informative prior; and an AR(1) model for the residuals.

3.2. PNC in Helsinki

Ultrafine particles are of interest in a range of areas, including physics (Hinds, 1999), urban planning (Moschandreas, 1998) and epidemiology (de Hartog et al., 2003; Krewski et al., 2009; HEI, 2010). In the urban environment it has been established that a significant portion of the ultrafine particles come from vehicles (Morawska et al., 2008; Virtanen et al., 2006; Pohjola et al., 2007; Hussein et al., 2007) and that the particle number concentration (PNC) varies non-linearly and irregularly over the day, week and with meteorology (Mejía et al., 2007; Morawska et al., 2002; Hussein et al., 2006; Ketzel et al., 2004; Pérez et al., 2010; Wu et al., 2008; Wehner and Wiedensohler, 2003).

Continuously measured time series of ultrafine PNC exhibit temporal trends, dependence on meteorology and autocorrelation (Wehner and Wiedensohler, 2003; Pérez et al., 2010; Wu et al., 2008; Hussein et al., 2004, 2006; Järvi et al., 2009). As such, the desire for flexible models which take these features into account without specifying the functional form of the relationship has motivated the use of the Generalised Additive Model with splines as basis functions and Generalised Linear Models with Fourier series basis functions.

In this case study we use hourly averaged size fractionated PNC, recorded at the SMEAR-III station at the Kumpula campus of the University of Helsinki using a twin DMPS system (Järvi et al., 2009). Meteorological data was recorded at the university campus on the rooftop of the Physicum building. For further information on the data collection, see sections 2.2-2.4 of Mølgaard et al. (2012).

Four different specifications of the model are fitted and model choice is made with the DIC. The reduction in autocorrelation is analysed and estimates of the fitted splines are provided.

The autocorrelation function of the ultrafine PNC in Helsinki is shown in Figure 4. This autocorrelation motivates the use of splines to estimate temporal trends and a model for the residuals which can capture any remaining variation.

Parameter(s)	Basis type	Basis size(s)	Periodic	Covariance structure
Hour of the day	B-spline	6	Y	2nd order penalty
Day of the week	Factor	7	N	iid
Day of the year	B-spline	6	Y	2nd order penalty
Wind direction	B-spline	8	Y	2nd order penalty
Wind speed	Thin plate	6	N	iid
Temperature	B-spline	8	N	2nd order penalty
Traffic	B-spline	8	N	2nd order penalty
Relative humidity	B-spline	8	N	2nd order penalty
Solar radiation	B-spline	8	N	2nd order penalty

Table 2: List of covariates and their univariate bases used in fitting a semi-parametric regression and forecasting model to the Helsinki data.

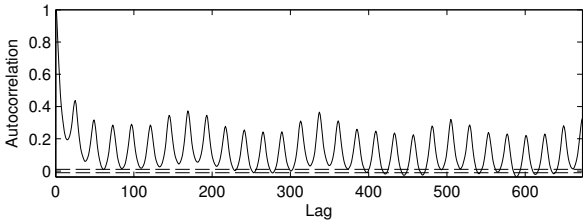

Figure 4: Autocorrelation function of ultrafine PNC in Helsinki for the data described in Section 3.2. Local maxima occurring at lags which are multiples of 24 correspond to a daily trend relationship, multiples of 168 to a weekly relationship. The maximum lag here corresponds to four weeks. The dashed lines represent the 5% significance level.

Table 2 describes the basis functions used in fitting the regression model. The thin plate spline used for wind speed includes a linear fixed effect and the random effects are five first order polynomials, $|x_i - \kappa_k|$. A tensor product of the wind speed and wind directions is used, with a basis size of 48. An attempt was made to use P-splines for the wind speed term but the presence of splines with no support at the extremes of the covariate space is undesirable.

Four models are fit and the "best" model is chosen with the DIC, calculated at the final stage of the model fitting. The first model is the one described by Table 2 and the joint wind speed and wind direction term described above. The second model replaces the univariate splines for temperature, traffic and relative humidity with a tensor product of each of those covariates with wind direction. The third model retains the univariate splines as well as the tensor products. The fourth model is the second model with the daily trend replaced with a tensor product of the daily trend and annual trend in order to recognise that the daily trend may change over the year.

The annual trend was excluded from the first three models as solar radiation and temperature exhibit very strong annual trends (FMI, 2011). Its inclusion in the fourth model is for the purposes of the tensor product with the daily trend. We will provide a plot of the fitted tensor as well as the marginal annual trend and marginal daily trend. These marginal trends are obtained by averaging over a prediction basis defined on a mesh consisting of one copy of each unique combination of time of day and day of the year.

For the autocorrelated residuals, lags at values of 1 hour, 24 hours and 168 hours are used. These lags represent, respectively, an attempt to capture the leftover hour to hour variation from fitting the daily trend, the leftover daily variation from fitting the daily trend and the leftover weekly variation after fitting the day of the week factor term.

The models will be fit to the first three years of the four years of data and predictions, including the autocorrelated residuals, will be made for the following year with the measured covariate values.

4. Results and Discussion

4.1. Simulation

MCMC estimation was conducted by drawing 5000 MCMC samples from the posterior; 500 initial samples are discarded as burn in.

We see (Figure 5) that the periodic term is estimated accurately by the model, as is the non-linear

interaction. We note that the credible intervals for the periodic term are quite similar in width across the entire covariate domain. This is due to the periodicity of the basis, uniform spacing of the covariate values and that each unique covariate value occurs the same number of times.

The fitted two dimensional spline captures the curvature and asymmetry of the simulation function (Figure 5).

Traces and densities of the posterior samples drawn from the posteriors of all parameters have converged and are unimodal and normally distributed except for the standard deviation parameter whose square is inverse Gamma (Figure 6).

Figure 7 gives trace and density plots for the smoothing parameters, which are Gamma distributed. We see that λ_x and λ_y have different, though not overly so, credible intervals. Had the 2D covariate been more oscillatory it is likely that we would end up with markedly distinct smoothing parameters (see Eilers and Marx, 2003, sec. 6). The priors for these smoothing parameters have a maximum at zero. The maximum in each of these posteriors now occurs at a non-zero value. While the parameters are still (approximately) Gamma distributed, they are no longer $\Gamma(1, \cdot)$.

The first parameter is the intercept term, β_0, parameters 2 to the 37 are the coefficients of the two dimensional spline, parameters 38 to 43 are the coefficients of the periodic spline. Parameter 44 is ϕ, the autoregressive parameter from the AR(1) model for the residuals and the final parameter is the standard deviation of the independent identically distributed errors, whose square is inverse gamma. All other parameters are marginally normally distributed.

The contours of the fitted 2D spline show that the simulated 2D covariate effect has been accurately reconstructed (Figure 8). The estimate for β_0 has been added to the 2D spline to allow for a more direct comparison of the values of the contours, as the fitted spline is centred about zero but the 2D covariate effect is not. Note that both the simulated temporal trend and its corresponding univariate spline are centred about zero and the AR errors also have a zero mean. The credible intervals for the two dimensional spline are quite wide in regions where there are not many observations and at the edges of the covariate space. The mean surface of the 2D smooth has no local maxima or minima, the presence of which would indicate excessive wiggliness.

Figure 9 shows that the residuals in $\boldsymbol{u}$ have

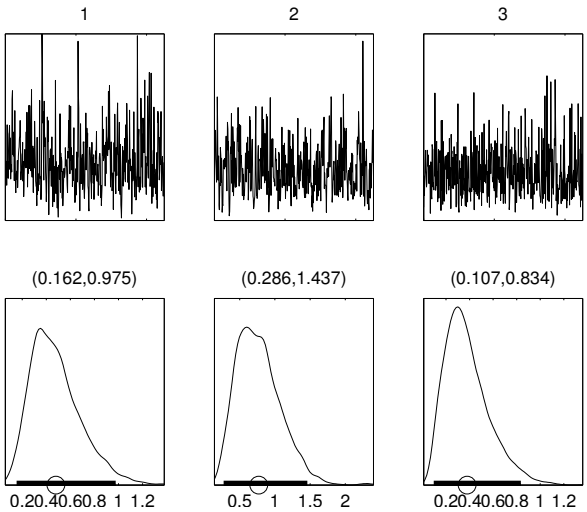

Figure 7: Trace and density plots for MCMC samples of smoothing parameters for simulated data. From left, λ_x, λ_y and λ_t.

smaller autocovariance than the sampled values of $\boldsymbol{\varepsilon}$. The diagonal bands correspond to 24 lags, indicating that the variation 24 observations apart may be modelled by expanding the model for the AR residuals to contain lags 1 and 24. Even so, the lag 24 bands and the background are lighter for $\boldsymbol{u}$ than for $\boldsymbol{\varepsilon}$ suggesting that modelling the autoregressive nature of the residuals has reduced the correlation of the posterior samples of the residuals.

The effective degrees of freedom of the intercept are concentrated around 1 with very little variance; it would be troubling were this not the case (Table 3). The 2D spline has between 21 and 27 effective degrees of freedom; the basis for this term has 36 elements, so each basis element requires, on average, fewer degrees freedom than would be required by the corresponding polynomial basis. Similarly, the size of the basis for the temporal trend is six basis elements (and would have been eight had we not insisted on periodicity of the basis) and the effective degrees of freedom of this term is slightly less than five.

4.2. PNC in Helsinki

All but the first year is forecast both one day and two days in advance. The forecast values are stored throughout the iterative forecasting such that all bar the first year of the data is forecast conditioned on the preceding data.

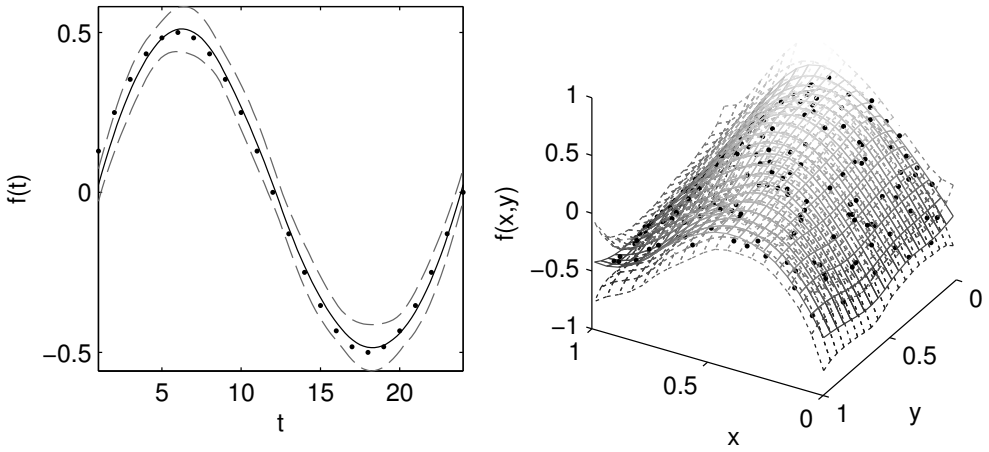

Figure 5: Marginal effects of splines (solid), 95% credible interval (dashed) and the true values from the simulated data (dots).

Term	2.5%	50%	97.5%	Mean	SD
Intercept	1.00	1.00	1.00	1.00	0.00
2D	23.34	25.61	27.45	25.57	1.03
Temporal	4.95	4.98	4.99	4.98	0.01
Total	29.32	31.59	33.43	31.55	1.03

Table 3: Effective degrees of freedom summary statistics for model fit to simulated data.

Of the four models fitted, the model with the joint annual-daily trend had the lowest DIC (Table 4). The triangle shape of the PIT is very similar for each of the models and indicates biasedness in the samples towards slight underprediction (Figure 10a). The estimates of the variance of the modelled values (Figure 10b) indicates that while the models provide estimates with similar variances, the joint annual-daily temporal trend with tensor product meteorology provides more concentrated forecasts. This model also has the lowest DIC out of the four models fitted. Therefore, this model represents the most efficient fit in terms of goodness of fit versus model complexity and has the most concentrated forecast values out of a group of models which perform similarly under the PIT. All subsequent analysis in this section is performed on this model.

Figure 11 shows the posterior density estimates and 95% credible intervals for the mean, standard deviation, (a) and (b), and the mean and 2.5% and 97.5% quantiles for the fitted smooth functions of temporal trends and non-linear covariate effects (c) – (i).

The daily trend in PNC varies over the year (see (c)) with most days having a peak around 10am. This trend peaks during the summer period (days 140–250) with a daily peak around 11pm and a plateau from approximately 6am to midday. The daily trend in PNC in winter has a trough at 3am and a peak around midday (when most of the day's light occurs). Spring and autumn daily trends contain two peaks, one in the late evening around 9pm and one in the morning aroung 9-10am. The overall shape of these temporal trends is consistent with previously reported temporal trends (Mølgaard et al., 2012).

The weekly trend (d) shows a decreased partial effect on weekends (days 1 and 7) and after accounting for these temporal trends, much of the remaining temporal variation is explained with the Lag 1 autoregressive error although there is still some amount of autocorrelation at Lags 24 and 168, which is captured by the model.

The joint effect of wind speed and wind direction (e) is roughly linearly decreasing and almost independent of wind direction for winds weaker than 4m/s. When the wind is blowing above 5 m/s there is an observable non-linear interaction of wind speed and wind direction such that at 11 m/s a wind of angle approximately 170 degrees (3 radians) corresponds to an increase in log PNC and at 30 degrees (0.5 radians) the contribution to log PNC is negative. That is, a strong wind can either remove

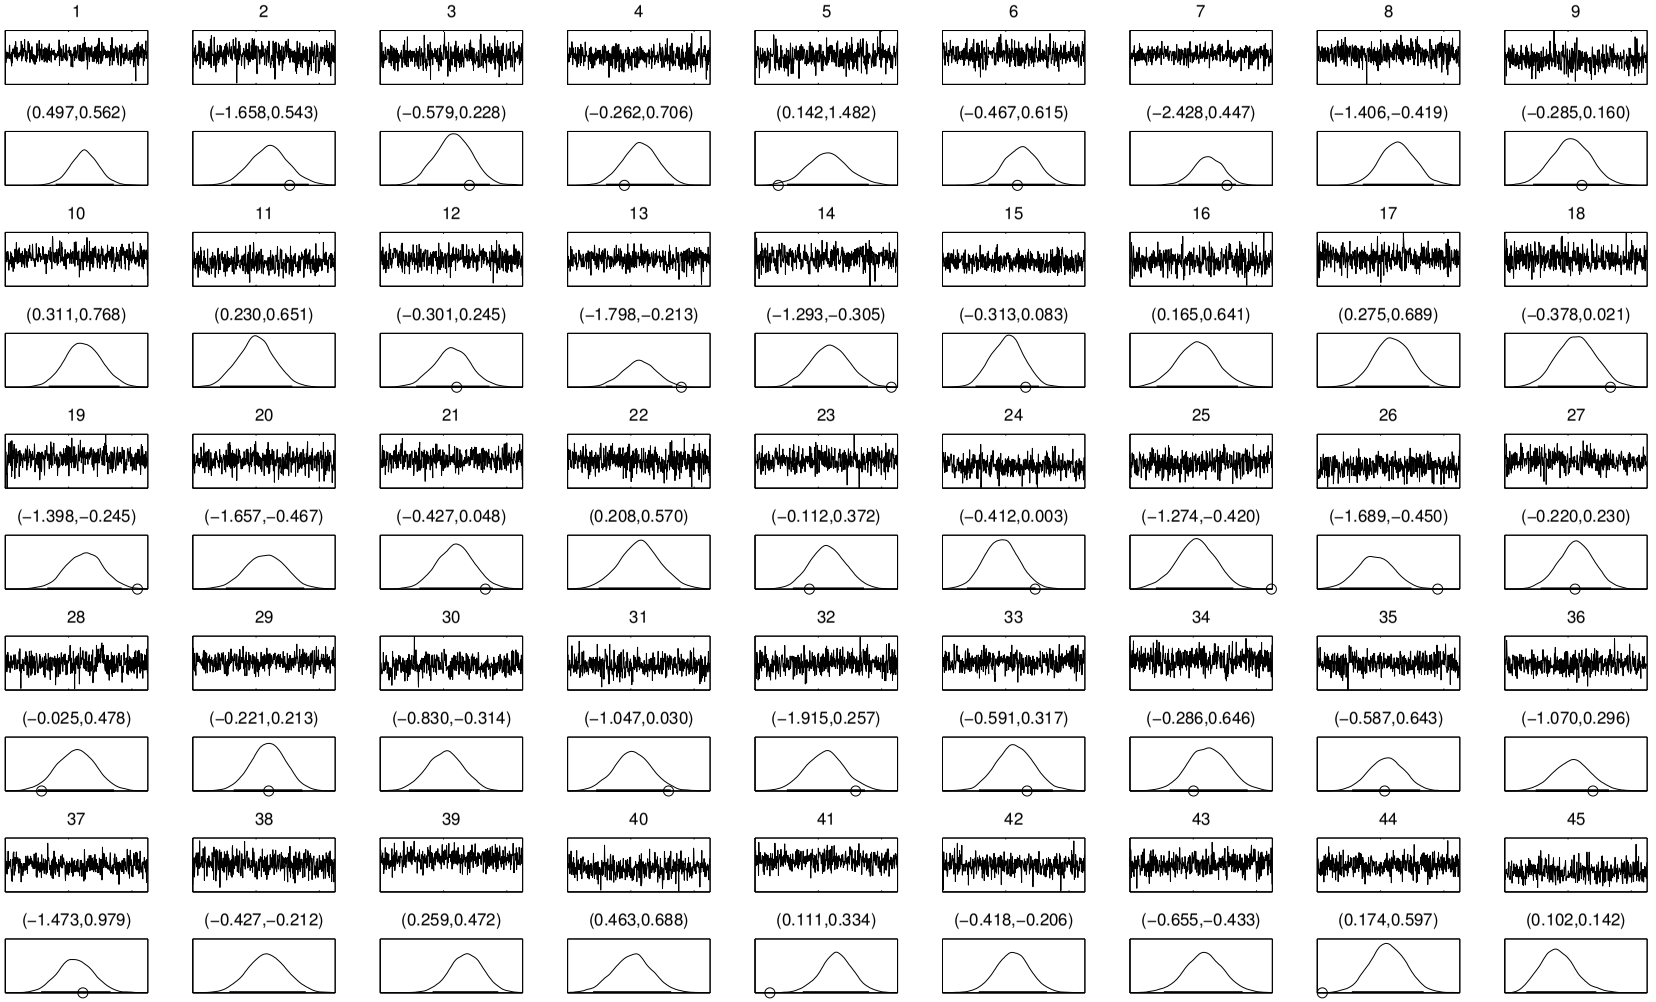

Figure 6: MCMC posterior summaries. The 95% credible interval is represented by a horizontal black bar in the density plot and the quantiles are given to three decimal places. Where zero is within the range of the posterior samples, it is plotted on the density plot as a circle.

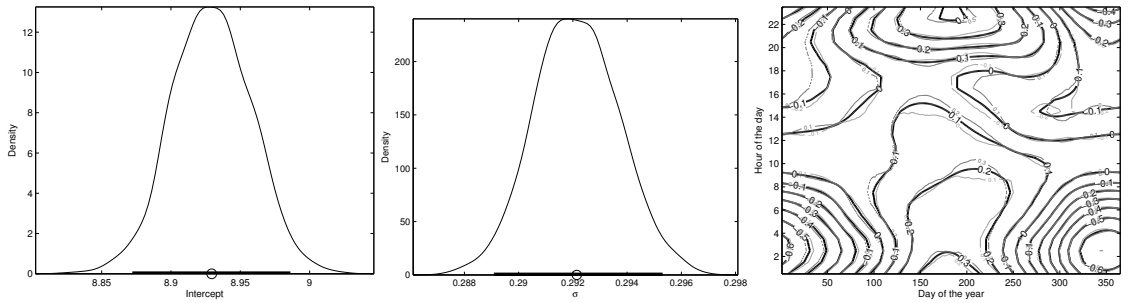

(a) Intercept (b) Standard deviation (c) Joint daily and annual trend

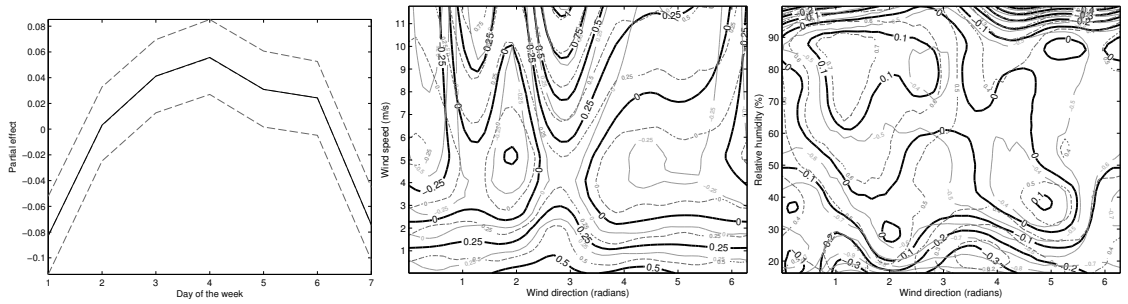

(d) Weekly trend (e) Joint wind effect (f) Humidity

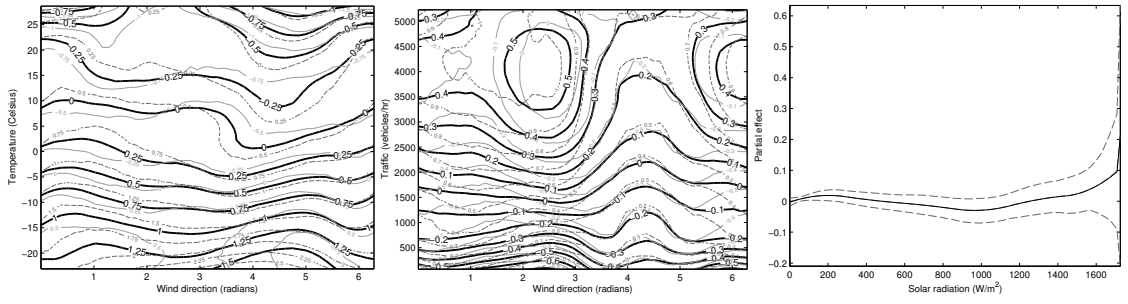

(g) Temperature (h) Traffic (i) Solar radiation

Figure 11: Density plots of β_0, σ, temporal trends and meteorological covariate effects. 95% credible intervals are represented as black rugplots for parameters and as dashed lines for fitted smooth terms. The solid grey lines in the contour plots represent the 2.5% quantiles and the dashed grey lines represent the 97.5% quantiles.

Model	DIC	p_D	$\overline{\text{edf}}$	rank(X)
Univariate meteorology	13023	96.809	86.7699	94
Tensor meteorology	12404	215.022	203.169	262
Combination of univariate and tensor	12406	218.360	205.405	286
Tensor meteorology with annual-daily trend	12055	244.166	230.318	292

Table 4: Summaries of fitting the models with differing treatments of meteorological covariates with respect to wind direction and temporal trends. DIC, effective number of parameters, mean effective degrees of freedom, and number of columns in the design matrix, X.

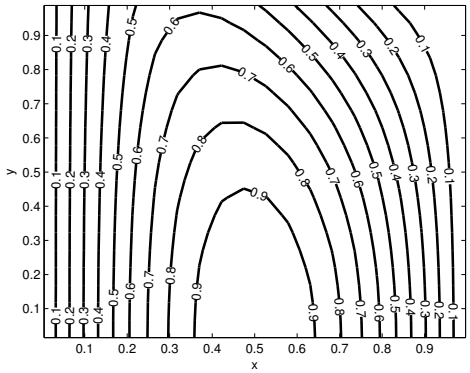

(a) True 2D covariate

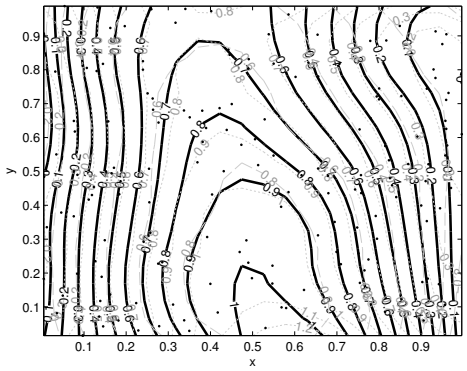

(b) Combined partial effects of 2D spline and constant

Figure 8: Comparison of contours of the simulated 2D covariate from 21 and the fitted 2D spline. Observed locations of x and y are plotted as points. The 2.5% quantile is represented with long grey dashes, the 97.5% quantile is represented with short grey dashes.

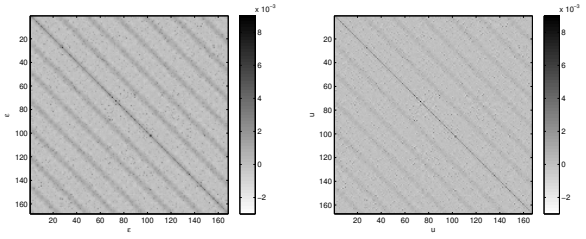

(a) Covariance of posterior samples of autocorrelated errors, ε_i, of each observation

(b) Covariance of posterior samples of iid errors, u_i, of each observation

Figure 9: Covariance matrices of posterior samples of autocorrelated errors, $\boldsymbol{\varepsilon}$ and independent identically distributed residuals $\boldsymbol{u}$, from simulated model.

particles from the microenvironment or transport them from a nearby source.

Marginalising this joint trend over the year provides an estimate of the mean daily trend in PNC and vice versa for marginalising over the day to obtain the mean annual trend (Figure 12. The mean daily trend in PNC exhibits two peaks, one at 10am and one at 10pm. The mean annual trend shows a maximum in summer, when the days are long.

Marginalising the posterior for the joint daily and annual trend (Figure 11), the mean daily trend in PNC exhibits two peaks, one at 10am and one at 10pm. The mean annual trend shows a maximum in summer, when the days are long (Figure 12).

The autocovariance of the posterior samples of the residuals for observations 1000 to 1200 are shown in Figure 13. The raw residuals, $\boldsymbol{\varepsilon}$, exhibit a noticeable amount of autocorrelation around observation 1050, indicating that while the semi-parametric regression model for the covariates has fitted a smooth annual and daily trend, not all the temporal variability has been explained. This residual variability is captured with by explicitly modelling the autocorrelation in the residuals and we

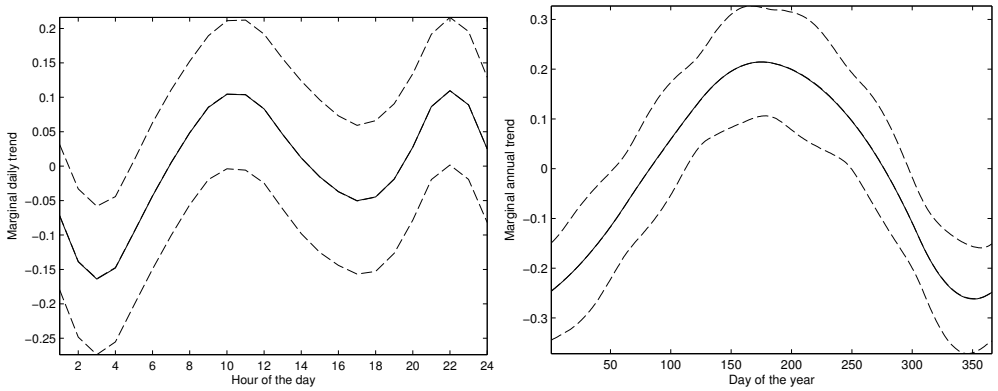

(a) Marginal daily trend (b) Marginal annual trend

Figure 12: Daily and annual trends calculated by marginalising the posteriors from the joint daily and annual trend.

see that the posterior covariance of these residuals decreases significantly when examining $\boldsymbol{u}$, the residuals with explicitly modelled autocorrelation.

The effect of high humidity (f) above 90%, is to reduce the total number of particles in the air because of precipitation. The effect in the middle range is fairly flat with some local peaks. The effect of temperature (g) is generally decreasing and shows some dependence on wind direction but there are no local peaks or troughs. The effect of traffic density (h) increases steadily up to about 3500 vehicles per hour with some dependence on wind direction. The increase is not as marked at a wind direction of about 230 degrees (4 radians), with a peak occurring around 140 degrees (2.5 radians).

Much of the effect of solar radiation 11i is accounted for by the joint daily and annual trend. The estimates of the effect of temperature, wind speed and direction, solar radiation and humidity (as a proxy for rainfall) and weekly trend show very good agreement with those previously reported (Clifford et al., 2011). In both the results presented here and by Clifford et al. (2011) the effects of temperature, humidity and wind speed are generally decreasing. Exceptions include strong wind speeds at certain headings and the positive effect of increased humidity until 90% as described above.

The estimates of $\boldsymbol{\phi}$ (Figure 14) do not contain zero in their credible intervals. We see a high amount of positive autocorrelation at Lag 1 and a small, but strictly positive, amount at Lags 24 and 168.

By specifying the regression model to include autoregressive residuals rather than autoregressive PNC, the mean daily trend (Figure 12) has peaks which occur at 10am and 10pm rather than at 8am and midday. The temporal variation in the model presented in this section can be expressed as

$$\log y_i = f\left(\text{hour}_i, \text{day of the year}_i\right) + \ldots + \left(1 - \phi_1 - \phi_{24} - \phi_{168}\right)^{-1} \varepsilon_i$$

while the temporal variation in the model of Clifford et al. (2011) corresponds to

$$\log y_i = f_1\left(\text{hour}_i\right) + f_2\left(\text{day of the year}_i\right) + f_3\left(\log y_{i-1}\right) + f_4\left(\log y_{i-24}\right) + \ldots + \varepsilon_i.$$

These two specifications are quite different as one contains a joint model of daily and annual trends and models the residuals autoregressively while the other models the residuals as autoregressive. As such, they provide quite different estimates of the daily trend. The weekly trend is the same across both models, with a maximum on Wednesday and minium values on the weekend.

To illustrate the modelling of the autoregressive nature of the residuals a contiguous subset of the modelled values and residuals was randomly selected, corresponding to observations 1000 to 1200. This is approximately eight days of measurements. Analysis of the autocovariance of the posterior samples of the residuals can be found in the appendix.

The effective degrees of freedom for the intercept, each spline term and the model overall are given in Table 5. Note that the intercept has exactly one

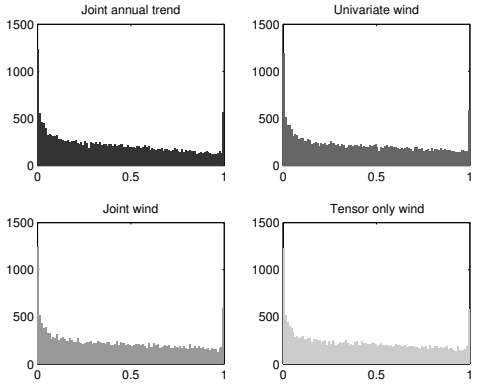

(a) Probability integral transforms

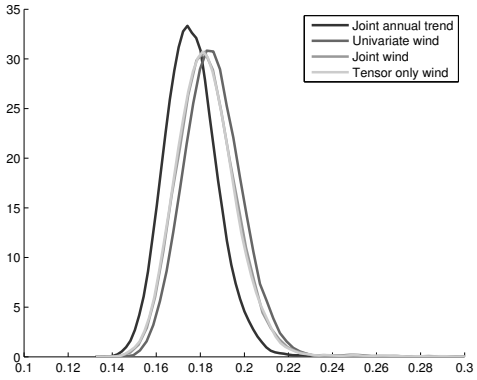

(b) Density estimates of the variance of the forecast values.

Figure 10: PITs and density estimates of forecast values for each of the four models fitted.

degree of freedom and that each spline has approximately one less effective degree of freedom than the number of basis splines; this is due to the constraint that each spline is centred around zero. The effective number of parameters for this model according to the DIC is 243.443. The effective number of degrees of freedom for the linear predictor is 230.643 (Table 5). These two estimates of model complexity agree quite well given the degrees of freedom given up by the constraint on each spline.

As described in Section 2.7, estimates of forecast values are obtained by modelling the observed values, modelling the future value based on measured meteorological measurements (in reality, weather forecasts are available instead of observed values) and adding the autocorrelated noise based on the estimates of $\boldsymbol{\phi}$. The modelled values (for both the observed and future values) are conditioned on $\boldsymbol{\phi}$ and so any inference performed on the forecasting

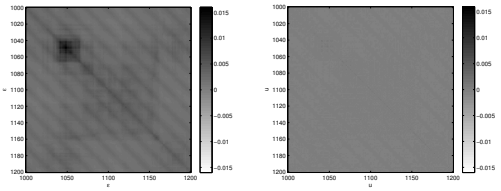

(a) Posterior covariance of $\boldsymbol{\varepsilon}$

(b) Posterior covariance of $\boldsymbol{u}$

Figure 13: Autocorrelation of residuals estimated from posterior samples of $\boldsymbol{\varepsilon}$ and $\boldsymbol{u}$.

must be done on $\mathrm{X}\boldsymbol{\beta}+\boldsymbol{\varepsilon}$ rather than the linear predictor $\mathrm{X}\boldsymbol{\beta}$.

5. Conclusion

This paper presents a regression method which combines semi-parametric regression, in the form of penalised splines, and a GLM with autoregressive residuals.

It was shown in the simulation case study of section 4.1 that the method described is capable of approximating the underlying smooth functions used to generate the data as well as the autoregressive noise which was added. The resulting smooth reproduced the underlying 2D function without the oscillations which characterise the use of B-splines with many basis functions and a small amount (e.g. none) of smoothing (Eilers and Marx, 2010).

In section 4.2 the modelling methodology was applied to some real data from Helsinki to infer the temporal trends and the effects of various meteorological and physical phenomenon on ultrafine PNC. The resulting fits were consistent with previous studies of ultrafine PNC in Helsinki (Clifford et al., 2011; Mølgaard et al., 2012) but provided fitted smooth functions which do not exhibit the oscillations typical of the use of a Fourier series basis.

Despite the models here having a basis size of more than 200 and a number of smoothing parameters and AR parameters, convergence is fast due to the use of the Gibbs sampler and the sum to zero constraint which ensures identifiability. The first 200 samples were discarded for each block of model fitting and the 2000 samples used for posterior inference came from stable chains, yielding normal posterior densities (e.g. Figure 11a).

Term	2.5%	50%	97.5%	Mean	SD
Intercept, β_0	1.00	1.00	1.00	1.00	0.00
Wind speed and direction	46.54	47.01	47.41	47.00	0.22
Solar radiation	6.93	6.97	6.99	6.97	0.01
Annual and daily trend	31.39	32.42	33.51	32.42	0.55
Weekly trend	6.00	6.00	6.00	6.00	0.00
Temperature and wind direction	38.14	40.55	43.28	40.59	1.37
Traffic and wind direction	48.30	50.12	51.99	50.15	0.96
Relative humidity and wind direction	44.22	46.44	48.92	46.52	1.18
Total for entire model	227.18	230.66	233.91	230.64	1.68

Table 5: Effective degrees of freedom summary statistics

By converting the univariate spline bases of the local meteorological covariates (traffic count, temperature, etc.) to a tensor basis of wind direction and the those covariates the DIC was reduced from 13023 to 12404, indicating that the effect of these covariates is dependent on wind direction. This analysis was absent from Clifford et al. (2011) and the move to a penalised B-spline basis improved the flexibility of the model over that described by Mølgaard et al. (2012). Fitting a model with the univariate bases and tensor bases for the covariates did not provide a qualitatively different fit and merely inflated the DIC.

The concurrent modelling of the residuals showed a high level of autocorrelation in the residuals at lag 1. The value of the autocorrelation parameters at lags 24 and 168 is small but non-zero and these lags correspond to modelling the variation in the residuals which is left after modelling the daily and weekly trends. Modelling this autocorrelation reduced the magnitude of the posterior covariance of the residuals.

Modelling with penalised splines captured the smooth temporal trends and the autoregressive model for the residuals explained the residual non-smooth variation. By combining these in the same model, rather than doing it in two steps, the MCMC sampler can trade the smooth trend and rough residuals off against each other; the posterior density for $\boldsymbol{\beta}$ contains $\mathrm{X}^* = (1 - \phi(L))\,\mathrm{X}$ and the posterior for $\boldsymbol{\phi}$ contains $\mathrm{E} = (\boldsymbol{y} - \mathrm{X}\boldsymbol{\beta})\left(1, L, L^2, \ldots, L^p\right)$. Fitting this model in a two-step process (i.e. modelling the autocorrelation in the residuals *post hoc*) would not allow this trade-off.

With respect to the methodology outlined in Section 2, the number of knots for each spline is fixed rather than allowing the number of knots in the splines to vary and using a reversible jump MCMC method (Biller, 1998). The advice of Eilers and Marx (2010) is to use a large number of knots, perhaps more than is "necessary", and to allow the smoothing penalty to control how wiggly the resulting smooth fit is. Ruppert et al. (2003) and Wand (2003) suggest choosing the number of knots for a univariate spline as $\min(n/4, 35)$ (for n the number of unique values the covariate takes) to ensure that the non-linear features are captured. The equivalent degrees of freedom for the spline will generally be substantially lower than the number of basis splines used. As such, p_D will typically not increase as more knots are used, because the DIC (and p_D) are calculated from the deviance (which stabilises as the basis size is increased).

Modelling with tensor products of splines allows the investigation of the interaction of two covariates which may have non-linear effects and non-linear interactions without specifying a particular functional form, which may often be no better than a subjective guess.

Adding random effects to a GLM converts it to a Generalised Linear Mixed Model. By analogy, random effects can be added to a GAM to form a Generalised Additive Mixed Model (GAMM) (Lin and Zhang, 1999). A mixed effect spline can be formed by taking the tensor product of a univariate (or higher dimension) spline with an identity matrix representing the different groups for the mixed effect. As the basis vectors stay the same across mixed effect groups, there is no need to include multiple copies of the basis. It is desirable to come up with a way to pass an argument to the model setup which allows the reuse of basis functions, similar to the `by=` argument in the R package `mgcv` Wood

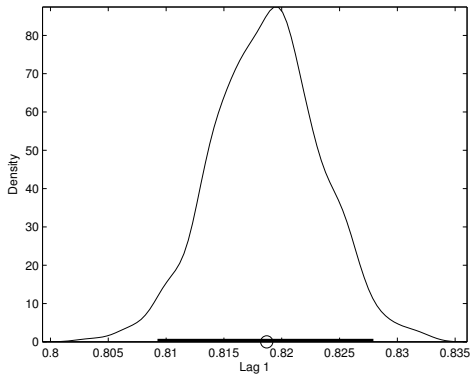

(a) Lag 1

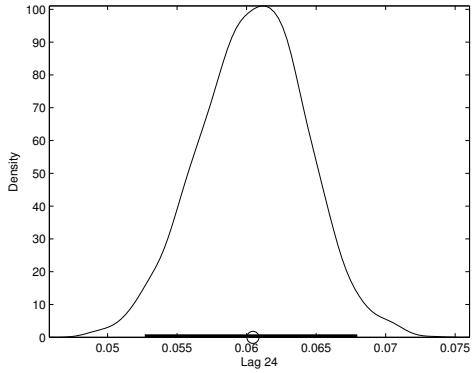

(b) Lag 24

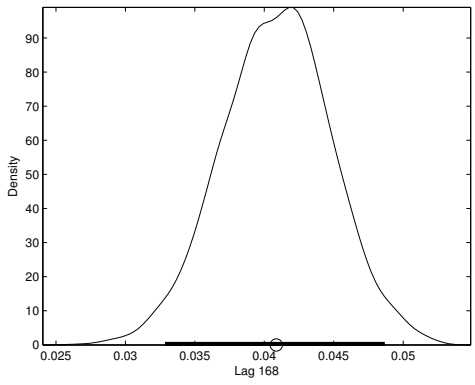

(c) Lag 168

Figure 14: Density plots of autoregression parameters $\boldsymbol{\phi}$. 95% credible intervals are represented as black rugplots.

(2011). This would be a computational improvement over forming the tensor product of a spline and factor term.

The method presented here provides flexible fitting of covariates which may have non-linear effects. The focus has been on temporal trends using cyclic B-spline basis functions with smoothing penalties to ensure there are no spurious oscillations. The anticipated autocorrelation of the residuals has been explicitly modelled to account for much of the temporal variation that remains after removing smooth trends.

Acknowledgement

The authors wish to thank the Division of Atmospheric Sciences and Faculty of Science at the University of Helsinki as well as the Institute for Health and Biomedical Innovation and the Discipline of Physics and School of Mathematical Sciences in the Science and Engineering Faculty at Queensland University of Technology, Brisbane, Australia for their support of Sam Clifford in travelling to Helsinki.

The authors also wish to thank Professor Markku Kulmala (University of Helsinki) and Professor Lidia Morawska (QUT) for their feedback and suggestions throughout the development of this work and Professor Tilmann Gneiting for his comments on a previous version of this paper.

References

Belitz, C., Brezger, A., Kneib, T., Lang, S., 2009. BayesX - Software for Bayesian inference in structured additive regression models. Version 2.0.1.

Biller, C., 1998. Adaptive bayesian regression splines in semiparametric generalized linear models. Journal of Computational and Graphical Statistics 9, 122–140.

Box, G.E.P., Jenkins, G.M., 1994. Time Series Analysis: forecasting and control. Prentice-Hall.

Chib, S., 1993. Bayes regression with autoregressive errors : A Gibbs sampling approach. Journal of Econometrics 58, 275–294.

Cleveland, R.B., Cleveland, W.S., McRae, J., Terpenning, I., 1990. STL: A seasonal-trend decomposition procedure based on loess. Journal of Official Statistics 6, 373.

Cleveland, W.S., Grosse, E., Shyu, W.M., 1993. Statisical Models in S. Chapman and Hall. chapter 8, Local Regression Models. pp. 309–376.

Clifford, S., Low Choy, S., Hussein, T., Mengersen, K., Morawska, L., 2011. Using the generalised additive model to model the particle number count of ultrafine particles. Atmospheric Environment 45, 5934–5945.

Crainiceanu, C.M., Ruppert, D., Wand, M.P., 2005. Bayesian analysis for penalized spline regression using WinBUGS. Journal of Statistical Software 14.

Dawid, A.P., 1984. Present position and potential developments: some personal views: statistical theory: the prequential approach. Journal of the Royal Statistical Society A 147, 278–292.

de Boor, C., 1977. Package for calculating with B-Splines. SIAM Journal on Numerical Analysis 14, 441–472.

de Boor, C., 1978. A practical guide to splines. Springer-Verlag.

Diebold, F.X., Gunther, T.A., Tay, A.S., 1998. Evaluating density forecasts with applications to financial risk management. International Economic Review 39, 863–883.

Eilers, P.H.C., Marx, B.D., 1996. Flexible smoothing with B-splines and penalties. Statistical Science 11, 89–121.

Eilers, P.H.C., Marx, B.D., 2003. Multivariate calibration with temperature interaction using two-dimensional penalized signal regression. Chemometrics and Intelligent Laboratory Systems 66, 159 – 174.

Eilers, P.H.C., Marx, B.D., 2010. Splines, knots, and penalties. Wiley Interdisciplinary Reviews: Computational Statistics 2, 637–653.

Findley, D.F., Monsell, B.C., Bell, W.R., Otto, M.C., Chen, B.C., 1998. New capabilities and methods of the X-12-ARIMA seasonal adjustment program. Journal of Business and Economic Statistics 16, 127–152.

FMI, 2011. Climatological statistics of Finland 1981-2010. Technical Report. Finnish Meteorological Institute.

Fong, Y., Rue, H., Wakefield, J., 2007. Bayesian inference for generalized linear mixed models. Biostatistics 8, 1–27.

Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B., 2004. Bayesian Data Analysis. Chapman and Hall.

Gneiting, T., Balabdaoui, F., Raftery, A.E., 2007. Probabilistic forecasts, calibration and sharpness. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 69, 243–268.

Hamill, T.M., 2001. Interpretation of rank histograms for verifying ensemble forecasts. Monthly Weather Review 129, 550–560.

de Hartog, J.J., Hoek, G., Peters, A., Timonen, K.L., Ibald-Mulli, A., Brunekreef, B., Heinrich, J., Tiittanen, P., van Wijnen, J.H., Kreyling, W., Kulmala, M., Pekkanen, J., 2003. Effects of fine and ultrafine particles on cardiorespiratory symptoms in elderly subjects with coronary heart disease. American Journal of Epidemiology 157, 613–623.

Harvey, A., Koopman, S.J., 1993. Forecasting hourly electricity demand using time-varying splines. Journal of the American Statistical Association 88, 1228–1236.

Hastie, T.J., Tibshirani, R.J., 1990. Generalized Additive Models. Chapman and Hall.

Hastings, W.K., 1970. Monte Carlo sampling methods using Markov chains and their applications. Biometrika 57, 97–109.

Hay, J.L., Pettitt, A.N., 2001. Bayesian analysis of a time series of counts with covariates - an application to the control of an infectious disease. Biostatistics 2, 433–444.

HEI, 2010. Traffic-related air pollution: a critical review of the literature on emissions, exposure and health effects. Special Report 17. Health Effects Institute, Boston, MA.

Heiler, S., Feng, Y., 2000. Data-driven decomposition of seasonal time series. Journal of Statistical Planning and Inference 91, 351–363.

Hinds, W., 1999. Aerosol technology: properties, behavior, and measurement of airborne particles. Wiley-Interscience, Wiley.

Hussein, T., Karppinen, A., Kukkonen, J., Härkönen, J., Aalto, P.P., Hämeri, K., Kerminen, V.M., Kulmala, M., 2006. Meteorological dependence of size-fractionated number concentrations of urban aerosol particles. Atmospheric Environment 40, 1427–1440.

Hussein, T., Kukkonen, J., Korhonen, H., Pohjola, M., Pirjola, L., Wraith, D., Härkönen, J., Teinilä, K., Koponen, I., Karppinen, A., Hillamo, R., Kulmala, M., 2007. Evaluation and modeling of the size fractionated aerosol particle number concentration measurements nearby a major road in Helsinki - Part II: Aerosol measurements within the SAPPHIRE project. Atmospheric Chemistry and Physics 7, 40814094.

Hussein, T., Puustinen, A., Aalto, P.P., Mäkelä, J., Hämeri, K., Kulmala, M., 2004. Urban aerosol number size distributions. Atmospheric Chemistry and Physics 4, 391–411.

Järvi, L., Hannuniemi, H., Hussein, T., Junninen, H., Aalto, P.P., Hillamo, R., Mäkelä, T., Keronen, P., Siivola, E., Vesala, T., Kulmala, M., 2009. The urban measurement station SMEAR III: Continuous monitoring of air pollution and surface-atmosphere interactions in Helsinki, Finland. Boreal Environmental Research 14, 86–109.

Ketzel, M., Wåhlin, P., Kristensson, A., Swietlicki, E., Berkowicz, R., Nielsen, O.J., Palmgren, F., 2004. Particle size distribution and particle mass measurements at urban, near-city and rural level in the Copenhagen area and Southern Sweden. Atmospheric Chemistry and Physics 4, 281–292.

Krewski, D., Jerrett, M., Burnett, R.T., Ma, R., Hughes, E., Shi, Y., Turner, M.C., III, C.A.P., Thurston, G., Calle, E.E., , Thun, M.J., 2009. Extended follow-up and spatial analysis of the American Cancer Society study linking particulate air pollution and mortality. Technical Report 140. Health Effects Institute. Boston, Massachusetts.

Lang, S., Brezger, A., 2004. Bayesian P-Splines. Journal of Computational and Graphical Statistics 13, 183–212. http://pubs.amstat.org/doi/pdf/10.1198/1061860043010.

Lenk, P.J., 1999. Bayesian inference for semiparametric regression using a Fourier representation. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 61, 863–879.

Lin, X., Zhang, D., 1999. Inference in generalized additive mixed models by using smoothing splines. Journal of the Royal Statistical Society B 61, 381–400.

Lunn, D., Thomas, A., Best, N., Spiegelhalter, D., 2000. WinBUGS - a Bayesian modelling framework: concepts, structure, and extensibility. Statistics and Computing 10, 325–337.

Marx, B.D., Eilers, P.H.C., 2005. Multidimensional penalized signal regression. Technometrics 47, 13–22.

Mejía, J.F., Wraith, D., Mengersen, K., Morawska, L., 2007. Trends in size classified particle number concentration in subtropical Brisbane, Australia, based on a 5 year study. Atmospheric Environment 41, 1064–1079.

Mølgaard, B., Hussein, T., Corander, J., Hämeri, K., 2012. Forecasting size-fractionated particle number concentrations in the urban atmosphere. Atmospheric Environment 46, 155–163.

Morawska, L., Jayaratne, E.R., Mengersen, K., Jamriska, M., Thomas, S., 2002. Differences in airborne particle and gaseous concentrations in urban air between weekdays and weekends. Atmospheric Environment 36, 4375–4383.

Morawska, L., Ristovski, Z., Jayaratne, E.R., Keogh, D.U., Ling, X., 2008. Ambient nano and ultrafine particles from motor vehicle emissions: characteristics, ambient processing and implications on human exposure. Atmospheric

Environment 42, 8113–8138.
Moschandreas, D. (Ed.), 1998. Design, Construction, and Operation of Healthy Buildings - Solutions to Global and Regional Concerns. American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.
O'Sullivan, F., 1988. Nonparametric estimation of relative risk using splines and cross-validation. SIAM Journal of Scientific and Statistical Computation 9, 531–542.
Pérez, N., Pey, J., Cusack, M., Reche, C., Querol, X., Alastuey, A., Viana, M., 2010. Variability of particle number, black carbon, and PM10, PM2.5, and PM1 levels and speciation: Influence of road traffic emissions on urban air quality. Aerosol Science and Technology 44, 487–499. 17th Transport and Air Pollution Symposium, Toulouse, FRANCE, JUN 02-04, 2009.
Pohjola, M., Pirjola, L., Karppinen, A., Härkönen, J., Korhonen, H., Hussein, T., Ketzel, M., Kukkonen, J., 2007. Evaluation and modelling of the size fractionated aerosol particle number concentration measurements nearby a major road in Helsinki - PART I: modelling results from the LIPIKA project. Atmospheric Chemistry and Physics 7, 4065–4080.
Rue, H., Martino, S., 2010. INLA – Functions which allow to perform a full Bayesian analysis of structured additive models using Integrated Nested Laplace Approximaxion. 0.0 edition.
Ruppert, D., Wand, M.P., Carroll, R.J., 2003. Semiparametric Regression. Cambridge University Press.
Ruppert, D., Wand, M.P., Carroll, R.J., 2009. Semiparametric regression during 2003-2007. Electronic Journal of Statistics 3, 1193–1256.
Schwarz, G., 1978. Estimating the dimension of a model. Annals of Statistics 6, 461–464.
Silverman, B.W., 1985. Some aspects of the spline smoothing approach to non-parametric regression curve fitting. Journal of the Royal Statistical Society B 47, 1–52.
Smith, M., Wong, C.M., Kohn, R., 1998. Additive nonparametric regression with autocorrelated errors. Journal of the Royal Statistical Society B 60, 311–331.
Spiegelhalter, D.J., Best, N.G., Carlin, B.P., van der Linde, A., 2002. Bayesian measures of model complexity and fit. Journal of the Royal Statistical Society B 64, 583–639.
Venables, W.N., Ripley, B.D., 2002. Modern Applied Statistics with S. Springer.
Virtanen, A., Rönkkö, T., Kannosto, J., Ristimäki, J., Mäkelä, J.M., Keskinen, J., Pakkanen, T., Hillamo, R., Pirjola, L., Hämeri, K., 2006. Winter and summer time size distributions and densities of traffic-related aerosol particles at a busy highway in Helsinki. Atmospheric Chemistry and Physics 6, 2411–2421.
Wahba, G., 1990. Spline Models for Observation Data. Society for Industrial and Applied Mathematics.
Wand, M.P., 2003. Smoothing and mixed models. Computational Statistics 18, 223–249.
Wand, M.P., Ormerod, J.T., 2008. On semiparametric regression with O'Sullivan penalized splines. Australian and New Zealand Journal of Statistics 50, 179–198.
Wehner, B., Wiedensohler, A., 2003. Long term measurements of submicrometer urban aerosols: statistical analysis for correlations with meteorological conditions and trace gases. Atmospheric Chemistry and Physics 3, 867–879.
Wood, S.N., 2003. Thin plate regression splines. Journal of the Royal Statistical Society B 65, 95–114.
Wood, S.N., 2006. Generalized Additive Models: An Introduction with R. Chapman and Hall/CRC.
Wood, S.N., 2011. GAMs with GCV/AIC/REML smoothness estimation and GAMMs by PQL. Comprehensive R Archive Network. 1.7-4 edition.
Wu, Z., Hu, M., Lin, P., Liu, S., Wehner, B., Wiedensohler, A., 2008. Particle number size distribution in the urban atmosphere of Beijing, China. Atmospheric Environment 42, 7967–7980.